Lexical Semantic Constraints

Dan Fass †

CSS/LCCR TR 89-11

†Advanced Systems Institute Visiting Post Doctoral Fellow

92-05700/

[To appear in James Pustejovsky (Ed.) *Semantics and the Lexicon*]

Lexical Semantic Constraints

Dan Fass
Centre for Systems Science
Simon Fraser University
Burnaby, British Columbia, Canada V5A 1S6.

Abstract

In linguistics and natural language processing, language has been described and theorized about in terms of constraints. Given that lexical semantics is the study of the meaning ('semantics') of words ('lexical'), the constraints that are the most natural focus for lexical semantics research are constraints on the meanings of words, i.e., what one might call *lexical semantic constraints*. A number of lexical semantic constraints are investigated: selectional restrictions (also known as preferences), assertions, relevance, salience, and restrictions on the internal semantic structure of nouns and verbs. The challenge for lexical semantics is to understand how such constraints are organized and used.

1. Introduction

Two trends can be identified in recent studies of the semantics of natural language. One trend is a growth of interest in lexical semantics. A second and so far independent trend is to view "semantic interpretation" or "understanding natural language" as the integration of diverse constraints from language and from context. I want to argue that this view of language as the integration of constraints is pertinent to lexical semantics and that a natural focus for lexical semantics research is semantic constraints possessed by words, that is, *lexical semantic constraints*.

In section 1, the two trends and the idea of lexical syntactic and semantic constraints are discussed. In section 2, selectional restrictions, probably the best known lexical semantic constraint, are presented and some of their deficiencies noted. Some less well known lexical semantic constraints are described in sections 3-5. In section 6, some conclusions are offered.

1.1. Lexical Semantics

In recent years, there has been a notable growth of interest in lexical semantics, with meetings (e.g., Testen, Mishra, & Drogo, 1984), reviews (e.g., Levin, 1985), books (e.g., Cruse, 1986), theories of lexical semantics (e.g., Pustejovsky, 1987), and applications of lexical semantics to text processing (e.g., Slator, 1988). Curiously though, a general definition of lexical semantics are lacking in this research, perhaps because such a definition has been thought self-evident. Nevertheless, and at the risk of stating the obvious, given the meanings of 'lexical' ("of words") and 'semantics' ("the study of meaning"), it follows that lexical semantics is a branch of semantics concerned with the study of the meaning of words (cf. Ullmann, 1957, p. 33) and other lexical items like idioms and fixed phrases. Topics for study therefore include lexical ambiguity, lexical decomposition, lexical structure, and the lexicon. Moreover, as suggested by the kinds of theories reviewed in Levin (1985), much work of the 60s and 70s might be included as lexical semantics research, for example, theories of lexical decomposition in what is often called "linguistic semantics" based on semantic features (e.g., Katz & Postal, 1964; Katz, 1972) or primitive actions (e.g., Jackendoff, 1976), sometimes emphasizing case or thematic roles (e.g., Gruber, 1965; Fillmore, 1968), plus computational theories of lexical decomposition in natural language processing (NLP) such as those based on semantic primitives (e.g., Schank, 1975; Wilks, 1975a), words (e.g., Hawkinson, 1975), word senses (e.g., Quillian, 1967; Fass, 1988a), and those emphasizing case or thematic roles (e.g., Charniak, 1981; Pustejovsky, 1987).

1.2. Semantic Interpretation as the Integration of Constraints

In recent years, a number of researchers in linguistics and NLP have described understanding natural language or semantic interpretation as the "use" or "integration" or "propagation" of constraints. Perhaps the clearest statement of this view has been made by Weischedel and Ramshaw (1987):

> "The key to understanding natural language is the use of constraints from language and from context to guide the search for the intended meaning by eliminating the plethora of alternative ways to group words. Those constraints make the problem tractable by reducing the size of the space that must be searched for an interpretation" (Ibid., p. 155).

A similar view has been expressed by Lehnert (1988), whose work in NLP has a connectionist orientation. She discusses syntactic and semantic constraints and their respective roles in semantic interpretation (cf. p. 3) and describes the architecture of a program called CIRCUS which has processing modules that contain

> "(1) local syntactic constraints, (2) local semantic constraints, and (3) global constraints of either a syntactic or semantic nature" (Ibid., p. 36).

The idea of propagating constraints is prominent in unification-based approaches to language understanding. For example, Fenstad, Halvorsen, Langholm, and Bentham (1987) outline

> "[a] framework for relating the *linguistic form* of utterances and their *semantic interpretation* which is based on the idea of constraint propagation" (p. 1, their italics).

Similarly, they describe their view of semantic interpretation as

> "the product of the integration of constraints from a multiplicity of levels of description of linguistic form" (Ibid., p. 12).

Later, they describe "a single pool of constraints" made up from morphology, syntax, and phonology, plus "contextual factors" (Ibid., p. 15). Interestingly, there is no mention of "semantic" constraints. Pollard and Sag (1987) describe another unification-based theory of language in which they acknowledge Fenstad et al's ideas about a pool of constraints (p. 8). Pustejovsky's (1987) theory of lexical semantics also includes a "constraint propagation method" (p. 172) for modelling concept acquisition.

It is also worth noting that a recent text book on computational linguistics (Grishman, 1986) is partially structured in terms of the specification and enforcement of syntactic and semantic constraints (see esp. pp. 111-121). Another recent text book on natural language understanding (Allen, 1987) mentions that

> "selection restrictions and case grammar indicate interword constraints" (Ibid., p. 218)

and the subheading of a later chapter is titled "Constraints Imposed by Syntax and Semantics" (p. 302) and is followed by a brief discussion of knowledge about the case structure and selection restrictions of verbs to resolve lexical ambiguity.

1.3. Lexical Syntactic and Semantic Constraints

Oddly enough, in none of the above research did I find the phrase "lexical semantic constraint" and nor for that matter "lexical syntactic constraint." This is strange because much of the work is lexicalist; that is, syntactic and semantic generalizations are located in the lexicon. This is so even in lexicalist syntactic theories such as Lexical-Functional Grammar (Kaplan & Bresnan, 1982), Word Grammar (Hudson, 1984), and Head-Driven Phrase Structure Grammar (Pollard & Sag, 1987).

The most relevant discussion concerning constraints that I came across is Starosta (1988), who emphasizes the importance of constraints in the formulation of syntactic theories (pp. 7-8) and distinguishes three kinds of constraint: constraints on grammatical representations, constraints on grammar rules, and – of the most interest here – constraints on lexical representations. Among those constraints on lexical representations, Starosta describes a series of "lexical features." Those features are divisible along two dimensions:

[1] contextual, non-contextual, or 'in-between';

[2] syntactic, semantic, or 'in-between'.

> "Contextual features specify ordering and dependency relationships among major syntactic categories ('parts of speech'), agreement and government requirements, case frames, and 'selection', semantic implications imposed by head items on their dependents. Non-contextual features characterize class memberships, including membership in purely syntactic categories such as [+N] and [+V], purely semantic, for example, [+dgre] 'degree', [+soft], etc., or in between, for example, [+mass], [+prnn] (pronoun), or [+plrl] (plural). The 'in between' features are semantic features which have external syntactic consequences, either in terms of entering into agreement relationships, such as [+plrl], or in restricting the range of elements allowed to cooccur in the same construction, such as [+mass] or [+prnn]. They include all inflectional features, including person, number, gender, and tense features as well as localistic case form and case relation features" (Starosta, 1988, p. 53).

The contextual/non-contextual distinction is implicit in most, if not all, syntactic theories. Starosta (1988, p. 52) points out that Chomsky (1965, p. 95) introduced separate syntactic and semantic contextual constraints: strict subcategorization and selectional restrictions. Fred Popowich (personal communication) points out that the same distinction applies to the LOCAL, SUBCAT, and NON-LOCAL features used in Head-Driven Phrase Structure Grammar reported in Pollard and Sag (1987) and to the separation of LOCAL and NONLOCAL information in a more recent version (Pollard, 1989). Table 1 summarizes and simplifies Starosta's observations, giving the following picture of lexical constraints from syntactic theory.

Table 1

Lexical constraints used in syntactic theory.

	Syntactic	Semantic
Contextual	Subcategorization	Selectional restrictions
Non-contextual	Syntactic features, e.g., tense, person, number, gender	Semantic features, e.g., [+soft], [-animate], [+human]

For Table 1, Starosta's observations have been simplified in two main respects. First, all of Starosta's contextual syntactic constraints have been classified as forms of subcategorization. Second, Starosta's 'in between' features have been included as non-contextual (since all constraints have some external consequences) and syntactic (since the features Starosta lists are quite often regarded that way).

2. Selectional Restrictions (Preferences)

Table 1 shows that the lexical semantic constraints identified in syntactic theory are selectional restrictions (which are contextual) and semantic features (which are non-contextual). Selectional restrictions are also known as selection restrictions (Katz, 1964) without the -al and have been much studied in linguistics (see, e.g., Lehrer, 1974, ch. 9). Katz describes selection restrictions as follows:

> "Each reading in the dictionary entry for a lexical item must contain a *selection restriction*, i.e., a formally expressed necessary and sufficient condition for that reading to combine with others. Thus, the selection restriction attached to a reading determines the combinations with the readings of other lexical items into which that reading can enter when a projection rule is applied" (Katz & Postal, 1964, p. 15, their italics).

A selectional restriction, then, is a restriction on the text surrounding a word sense. It is a "co-occurrence" constraint which is possessed by senses of words of certain parts of speech and which restricts the semantic classes of words with which it co-occurs, hence an adjective sense has a restriction upon the semantic class of nouns with which it co-occurs, and a verb sense has restrictions upon the semantic classes of nouns that fill its case roles. For example, the main sense of the verb 'drink' has the restriction that an

animal fills its agent case role, i.e., it is animals that drink.

Selectional restrictions have also been used extensively in NLP (see Hayes, 1976), where they have been called preferences (Wilks, 1975a) and expectations (Schank, 1975). This one constraint has been given several names because it has been put to different uses. In Katz's use, selection restrictions may or may not be satisfied, but when not satisfied (in other words: violated), they are assumed to be ill-formed and therefore left uninterpreted. Wilks uses the word 'preference' because he emphasizes that the semantic constraints require interpretation, whether they are satisfied or not. For further discussion (see Fass & Wilks, 1983). Schank stresses that the semantic constraints are used for top-down prediction, hence the term "expectation."

Preferences have been studied in a number of roles. One role is in lexical definition, to help discriminate the meanings of words. In this role, lexical semantic constraints serve to limit the semantic structure of words and their senses by specifying the similarities and differences between those words and their senses with respect to the constraints. To give an example, selection restrictions can help distinguish among senses of the verb "to play." For instance, the senses "to play a musical instrument" and "to play a game" can be distinguished, at least in part, because the senses have different selection restrictions as an object: a musical instrument and a game.

Two others roles are both in semantic interpretation. One of those roles is in distinguishing literal readings of text from anomalous or metaphorical ones (e.g., Wilks, 1975b). A second role is in lexical ambiguity resolution, as in Katz and Fodor (1963), Wilks (1975a) and Hirst (1987).

(1) "Mary plays the guitar."

For example, give the above two senses of 'play' and sentence (1), the differing preferences of the two senses can help select the intended sense,"to play a musical instrument."

Preferences, however, have a number of limitations. Selection restrictions work best in discriminating homonymous senses of words (cf. Olawsky, 1989, p. 171), where the meanings are unrelated and therefore distinct, rather than polysemous senses, where the meanings are related and thus often overlap. Hence, it is no coincidence that the verbs and nouns selected for many studies of lexical ambiguity resolution in linguistics and NLP have either homonymous senses or distinct polysemous ones. For example, Wilks (1975a, p. 59) uses two polysemous but distinct senses of 'grasp', one meaning "grasp a physical object," the other meaning "grasp an idea." In my own work, I have used twelve senses of the verb 'play' because their meanings were sufficiently distinct (there were 29 senses in the dictionary from which I drew those twelve). Another verb used is 'kick' which has the two homonymous senses meaning "kick a physical object (ball)" and "kick a habit."

The same careful selection of examples is true for nouns (though in these cases the deficiency is not preferences but rather the granularity of the semantic features used). For example, Katz and Fodor (1964, p. 507) use senses of 'ball' having the distinct meanings "a social activity" and "physical object having globular shape," though they complicated matters for themselves by introducing a third sense closely related to the second, meaning "solid missile for projection by engine of war." Another popular pair of noun senses is 'crook' meaning "a criminal" and "a shepherd's stick" (see, e.g., Wilks, 1975b, p. 334).

The careful selection of examples extends to treatments of prepositional ambiguity resolution. Prepositions have preferences for the subject and object they take (see, e.g., Wilks, Huang, & Fass, 1985). Examples of senses of a preposition must be chosen where the preferences of those senses are distinct; for example, two senses of 'with' that can be readily distinguished are the sense meaning accompaniment, which takes an animate object (e.g., "the person with the dog"), and the sense meaning instrument (e.g., "the person with the hammer").

A second problem with preferences is that they are sometimes very general and hence do not offer much in the way of constraints. This is true of certain common verbs and prepositions. For example, general verbs like 'be' and 'have' have virtually no preference restrictions on the subject and object they take. Likewise with (spatial) prepositions like 'in' or 'on' where the subject and object they take (in isolation) can be virtually any physical object. In such cases, the strongest constraint seems to be *relational*; i.e., a constraint that specifies some kind of restricted connection between subject and object.

(2) "The box was in the pen" (Bar-Hillel, 1960, p. 158).

For example, to resolve the lexical ambiguity of 'pen' in (2) requires enforcing the containment constraint implied by 'in' that the pen contains the box, hence the pen must be larger than the box, and so on.

(3) "The baby picked up the pen."

Hayes (1976) discusses two particular kinds of relational constraint. One kind he calls "physical constraints." In (3), Hayes observes that that if 'pen' can mean either a writing implement or a baby's play pen then in the above sentence 'pen'

> "can only refer to a writing implement, because it is beyond the strength of a normal baby to pick up a properly designed play pen .... In this example, something analogous to a selectional restriction is being used. The main difference is that the restriction involves a relation between the subject and object of "pick up", viz. that the physical strength of the subject must be adequate to cope with the weight of the object. In all of the selectional restrictions we considered before, the restrictions on the subject and object were independent" (Ibid., p. 54).

(4) "The gardener picked some stock."

(5) "The rancher picked some stock."

A second kind of relational constraint Hayes discusses is *topic*. Hayes contrasts the senses of 'stock' involved in (4) and (5). In the first sentence,

> "the sense of "stock" being used would be taken as that of a kind of sweet smelling plant, presumably because of the association between plants and the profession of gardener" (Ibid., p. 53),

whereas in the second,

> " "stock" would be taken to refer to cattle because of the parallel association between cattle and the profession of rancher" (Ibid.).

I would argue that spreading activation approaches (e.g., Quillian, 1967; Hirst, 1987) and connectionist approaches (e.g., Waltz & Pollack, 1985) exploit relational constraints in finding connections between subject and object.

(6) "Nadia's plane taxied to the terminal" (Hirst, 1987, p. 100).

For example, in (6), Hirst supposes that 'plane' can mean either "aeroplane" or "woodworker's tool" and that 'terminal' can mean either "computer terminal" or "airport building." In Hirst's approach to lexical ambiguity resolution, (6) is resolved in part finding a topical constraint between subject and object: a path is found between the aircraft-related senses of 'plane' and 'terminal'.

(7) "The astronomer married the star."

Waltz and Pollack (1985) present an interesting connectionist analysis of (7) in which a topical constraint between subject and object compete with the selectional restriction on the object of 'marry' and *interfere* with the correct resolution of (7) by favoring the celestial sense of 'star' (which is topically related to 'astronomer').

A third problem with preferences is that they fail to capture certain differences of meaning between words which have identical preferences. This is notable in modifiers such as adjectives and adverbs. For example, the adjectives 'male' and 'female' have the same preference for an organism (i.e., an animate thing like an animal or plant), yet they do not mean the same thing. What part of their meaning are preferences not accounting for? The same deficiency occurs with adverbs such as 'quickly' and 'slowly' which both apply to movement, but state different things about movement.

3. Assertions

To overcome the third problem requires the use of a second but less frequently studied type of lexical semantic constraint, which I call the "assertion." The notion of assertions is not new. The assertion of semantic information was noted by Lees (1960) in the formation of noun phrases and later developed by Katz and Fodor (1964) as part of the process of "attribution."

> "Attribution is the process of creating a new semantic unit compounded from a modifier and a head, whose semantic properties are those of the head, *except that the meaning of the compound is made more determinate than that of the head by the information the compound obtains from the modifier*" (Ibid., p. 508, italics in original).

Katz and Fodor gives some examples of attribution.

> "The word *aunt* is indeterminate as to age (i.e., neither the sentence "My aunt is an infant" nor "My aunt is aged" has any special semantic properties), but *spinster* ... contains the semantic marker *(Adult)* interpretation for *spinster aunt*. Thus, *spinster aunt* is made more determinate (with respect to age) than is *aunt*. This shows up in a comparison between the sentences "My spinster aunt is an infant" and "My spinster aunt is aged," the former of which is contradictory while the latter is not" (Ibid., p. 509).

Katz and Fodor note that "the limiting case" is "where the addition to the compound of semantic material from the modifier is zero," as in the compound *unmarried bachelor*.

> "The failure of the modifier to add semantic information would appear to account for the intuition that such expressions as *unmarried bachelor* are redundant and that, correspondingly, such statements as "Bachelors are unmarried" are *empty, tautological, vacuous, uninformative*" (Ibid., italics in original).

Like preferences, assertions are co-occurrence constraints; i.e., they are constraints upon the language that can co-occur with a word sense. The main difference between the two kinds of constraint is their *function*: preferences specify how the local context (of a word) be *restricted* whereas assertions specify information to be *added* to that context. An analogy may help clarify this functional difference: a preference is like the guidance system of a guided missile, which specifies (i.e., restricts) the missile's target, while an assertion is like the warhead, which is exploded on (i.e., is "added" to) the target.

Assertions are possessed by modifiers and are imposed on the heads to which the modifiers apply. For example, the adjective 'unmarried' contains information that any noun over which it predicates has unmarried marital status, hence "an unmarried giraffe" is a giraffe that is not married.

It is worth noting that modifiers possess both preferences and assertions, so they have *two* co-occurrence constraints. For example, 'female' has a preference for an organism and an assertion that the head it modifies be of the female sex. Assertions appear to play a significant role in defining the meaning of modifiers. For example, as noted in the last section, the meanings of the adjectives 'female' and 'male' are not distinguished by their preferences; rather, the difference in meaning is due (largely) to their differing assertions.

Another difference between preferences and assertions is the different *effects* that they produce when evaluated. I have argued elsewhere (e.g., Fass, 1988a) that preferences are the basis for four *preference-based semantic relations* (literal, metonymic, metaphorical, and anomalous) which in turn are the basis of literalness, metonymy, metaphor, and anomaly, and that assertions are the basis of three other *assertion-based semantic relations* (novel, redundant, and inconsistent) on which novelty, redundancy, and inconsistency are based.

(8) "Colorless green ideas sleep furiously" (Chomsky, 1957, p. 15).

To give an example of how the two semantic constraints can be used, consider (8), Chomsky's famous "nonsensical" sentence. To start with, 'colorless', 'green' and 'furiously' are all modifiers that each contain a preference and an assertion. The details of the preferences and assertions are unimportant here; what is worth emphasizing is that (8) contains seven contextual lexical semantic constraints (four preferences, three assertions) and that the overall nonsensical effect results from evaluating all seven constraints.

Let us begin by examining "green ideas," assuming that the sense of green meant here is the colour sense (rather than, for example the sense meaning "unripe"). That sense of 'green' has a preference; the effect produced by evaluating it with 'ideas' is anomaly. The same sense also has an assertion which produces an effect of novelty – after all, "ideas that are green" is a novel idea. Now let us consider the application of 'colorless' to "green ideas." The preference of 'colorless' (presumably the same as for 'green') applied to the adjective-noun produces a second effect of anomaly. The assertion of 'colorless' produces

an effect of inconsistency because an idea cannot be both green and colourless. Moving on to 'sleep' and 'ideas', 'sleep' has a preference that produces an effect of anomaly. And finally, with 'sleep' and 'furiously', my intuitions are that the effects produced are literalness and (mild) inconsistency.

So, the nonsensical effect of (8) can be decomposed into one instance of literalness, three instances of anomaly, one case of novelty, and two instances of inconsistency, with the cases of anomaly and inconsistency producing much, if not all, of the nonsensical effect. Readers may disagree with my intuitions and maybe also with the order in which I analyzed the parts of (8) (for example, perhaps the first analysis should have been 'colorless' and 'green') but my main concern is not to defend a particular interpretation of (8) but to repeat that any interpretation of (8) consists of evaluating a series of contextual lexical semantic constraints of (at least) two different kinds, preferences and assertions. As can be seen, the merit of using assertions is that they allow a richer linguistic description and that assertion-based semantic relations, in tandem with preference-based ones, allow us to build a multi-faceted semantic interpretation of a sentence.

Let us now look at how assertions might be represented in a lexical semantic theory and how the three assertion-based semantic relations can be discriminated. To do this, we will use Collative Semantics (hereafter CS), a computational lexical semantic theory (see, e.g., Fass, 1988b). For a description of the treatment of preference-based semantic relations, see Fass (1988c).

CS developed from an examination of Yorick Wilks' (1975b, 1978) Preference Semantics. Some early observations are reported in Fass (1983) and Fass and Wilks (1983). CS consists of four components. The first is the *sense-frame* which represents the meanings of word senses. The sense-frame definition of a word sense consists of other word senses defined in terms of other word senses, and so on, like in Quillian's (1967) memory model. The second component is *collation* which matches together pairs of sense-frames and thereby discriminates the semantic relations between the word senses represented by the sense-frames. The third component is the *semantic vector* which represents the matches found by collation and hence the semantic relations discriminated by that process. The fourth component is *screening* which finds the strength of relationship between two word senses by examining the semantic vector produced for the match between them; such relationships are compared in order to resolve lexical ambiguity.

CS has been implemented in a computer program called meta5 (so-called because it does more than metaphor) which analyzes short sentences containing lexical ambiguity. Meta5 discriminates the seven semantic relations between pairs of word senses in those sentences and resolves the lexical ambiguity in them. The assertion-based semantic relations can each be divided into two subkinds, making six subkinds of relation in all: redundant (overgeneral), redundant (identical), inconsistent (contradictory), inconsistent (contrary), novel (specific), and novel (unknown).

A redundant relation occurs when asserted information in a modifier is:

- *more general* than information in the head (redundant-overgeneral),
- *identical* to information in the head (redundant-identical).

(9) "John McEnroe is male."

A redundant-identical relation occurs between the adjective sense male1 and the proper noun john_mcenroe in (9).

```
sf(male1,                          sf(john_mcenroe,
   [[arcs,                            [[arcs,
     [[superproperty, sex1],            [[superinstance, tennis_player1]]],
      [property, male1]]],              [node0,
    [node1,                              [[sex1, male1],
     [[preference, organism1],           [marital_status1, married1]]]]).
      [assertion,
         [[sex1, male1]]]]]]).
```

Figure 1. Sense-frames of male1 and john_mcenroe.

Figure 1 shows the sense-frames for male1 and john_mcenroe. Sense-frames are made from two- and three-element lists called cells. All sense-frames contain two parts, an arcs section and node section. The arcs section contains genus information about the word sense described in the sense-frame. The information is represented as labelled arcs to superordinates, synonyms and antonyms; together, those arcs comprise a hierarchically organized network called the "sense-network." The node section contains differentia information. The node sections of nouns consist of lists of attributes concerning their structure (e.g., their composition and spatial extent) and function (e.g., who uses them). Only the most specific information about a noun sense is stored in its sense-frame; more general information is inherited down the sense-network when required by the process of collation. The information for all other parts of speech is subdivided into preference information and assertion information; in other words, the two co-occurrence constraints structure that information.

Asserted cell of male1	Cells of john_mcenroe
[sex1, male1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [composition1, flesh1], [animacy1, living1], [john_mcenroe, drink1, drink1], [john_mcenroe, eat1, food1], [john_mcenroe, be1, warm_blooded1], [biology1, human_being1], [john_mcenroe, have1, [2, leg1]], [john_mcenroe, have1, [2, arm1]], [john_mcenroe, locomote1, [on3, land1]], [john_mcenroe, play5, tennis1], **[sex1, male1],** [marital_status1, married1]]

Figure 2. Match of asserted cell from male1 against cells of john_mcenroe.

Figure 2 shows all the information about john_mcenroe inherited down the sense-network. The collation process matches [sex1, male1], the assertion of male1, against that information. In collation, matches of cells are classified as one of seven different kinds (ancestor, same, sister, descendant, estranged, unmatched source cells, unmatched target cells). The one found in Figure 2 is a *same* cell match, which indicates the redundant-identical relation.

An inconsistent semantic relation is one in which asserted information is *incompatible* with information in the head. Though CS and meta5 can discriminate inconsistent relations, neither can tell apart contradictory and contrary relations, the two subkinds.

(10) "John McEnroe is female."

There is an inconsistent relation between the adjective sense female1 and john_mcenroe in (10). The sense-frame for female1 is identical to the one for male1, except that it contains a different antonym (male1 instead of female1) and assertion.

Asserted cell of female1	Cells of john_mcenroe
[sex1, female1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [composition1, flesh1], [animacy1, living1], [john_mcenroe, drink1, drink1], [john_mcenroe, eat1, food1], [john_mcenroe, be1, warm_blooded1], [biology1, human_being1], [john_mcenroe, have1, [2, leg1]], [john_mcenroe, have1, [2, arm1]], [john_mcenroe, locomote1, [on3, land1]], [john_mcenroe, play5, tennis1], **[sex1, male1],** [marital_status1, married1]]

Figure 3. Match of asserted cell from female1 against cells of john_mcenroe.

Collation compares the assertion from female1 with the information about John McEnroe inherited down the sense-network, as Figure 3 shows. Of the seven kinds of cell matches, the one found in Figure 3 is called a *sister* cell match because of the relationship found between female1 and male1 in the two cells. As female1 and male1 are word senses, they have their own sense-frames which belong to the sense-network so a sense-network can be found between them, which is shown in Figure 4.

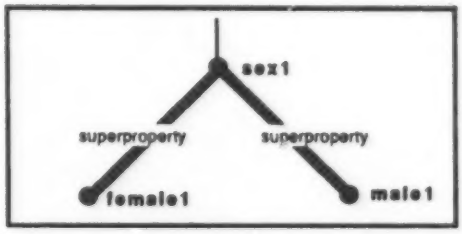

Figure 4. Sense-network path between female1 and male1.

Collation finds network paths using a graph search algorithm that searches for five different kinds of path (called: ancestor, same, sister, descendant, estranged). The one found in Figure 4 is called a sister path, hence the name given to the cell match. The sister cell match indicates an incompatibility between the cells, hence the inconsistent relation.

In a novel relation, the asserted information in a modifier is:

- *more specific* than information in the head (novel-specific), or
- *unlike* any pre-existing information in the head (novel-unknown).

(11) "The stick is oak."

An example of a novel-specific relation occurs between oak1 and stick1 in (11). Figure 5 shows the sense-frames for the two word senses.

```
sf(oak1,                                    sf(stick1,
   [[arcs,                                     [[arcs,
     [[superproperty, wooden1],                  [[superform, tool1]]],
      [property, oak1]]],                       [node0,
    [node1,                                      [[composition1, wood1],
     [[preference, physical_object1],             [human_being1, use1, it1]]]])
      [assertion,
        [[composition1, oak1]]]]]]).
```

Figure 5. Sense-frames of oak1 and stick1.

The assertion from oak1 is compared against the information for stick1 inherited down the sense-network, shown in Figure 6, and matches with [composition1, wood1].

Asserted cell of oak1	Cells of stick1
[composition1, oak1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], **[composition1, wood1],** [human_being1, use1, stick1]]

Figure 6. Match of asserted cell from oak1 against cells of stick1.

The match is called a *descendant* cell match because of the network path found between oak1 and wood1, which is a descendant path, shown in Figure 7. The discovery of this path indicates that the asserted cell is the more specific information and hence that the semantic relation here is novel-specific.

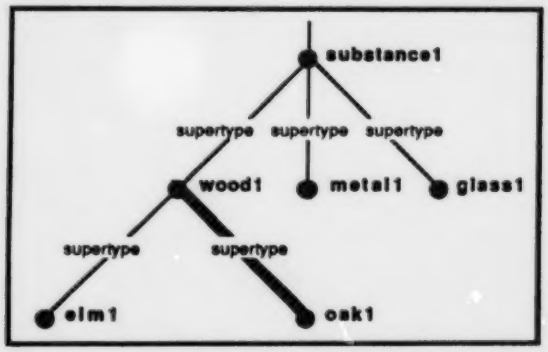

Figure 7. Sense-network path between oak1 and wood1.

Incidentally, the reverse network path is called the ancestor path and is used in an *ancestor* cell match, which indicates that the assertion is the more general, and hence is used to recognize a redundant-overgeneral relation, as in (12). The meta5 program analyzes (12) as well as sentences (9) through (11).

(12) "The stick is physical."

The key, then, to discriminating the assertion-based semantic relations is determining how asserted information matches against pre-existing information in the head. In CS, this is done by matching the information in terms of network paths between the word senses that represent the information. The algorithm in collation is described in Fass (1988b, 1988d). The very same network paths and cell matches are used for discriminating preference-based semantic relations as well.

4. Relevance and Salience

Relevance and salience are two constraints which both "highlight" or "prioritize" lexical semantic information. A form of relevance will be distinguished that is lexically based and uses such a restricted notion of context that, arguably, it acts like a non-contextual lexical semantic constraint. A version of salience will be described which prioritizes information inherited down a taxonomic hierarchy. CS will be used to demonstrate these notions.

Ortony, Vondruska, Foss, and Jones (1985) distinguish between relevance and salience as follows:

> "Relevance is a characteristic of attributes with respect to task and object domains .... Salience ... refers ... to the prominence or importance of an attribute in a person's representation of an entity or category" (p. 570).

The following examples may help clarify the distinction. A salient feature of many people's representations of foxes is that they are sly, unlike say gorillas which are strong, and so on. This is an unchanging attribute of foxes regardless of the context in which foxes are discussed. However, what is relevant about foxes varies with context: hence, if the context is some discussion of animal eating habits then what is relevant about foxes is their eating habits, and if the context is animal fur colouring then what is relevant about foxes is their fur colour. Relevance is thus a dynamic, domain-dependent notion: when the topic changes then so does what is relevant and hence those attributes that are relevant. Salience, on the other hand, is domain-independent and relatively fixed. Attributes that are salient remain so regardless of changes of domain. Salience is pretty clearly a non-contextual lexical semantic constraint. The status of relevance is less clear. At first glance, relevance looks like a general contextual notion and not obviously a lexical one, a constraint which is a property of the domain or background or setting of a piece of language.

In CS, however, a simple lexical notion of relevance is used: what is relevant in a sentence is the sense of the main sentence verb (the verb must be reasonably specific like 'drink' or 'plough' in order to select just a single relevant feature or attribute; a general verb like 'be' or 'have' selects several relevant features each). Like the description of relevance given by Ortony et al, this notion is dynamic and domain-dependent: the verb sense under examination indicates the current relevant domain and if the verb sense changes then so does what is relevant.

Relevance is used in CS in metonymy and metaphor interpretation. It is central to the view of metaphor taken in CS, that a metaphor consists of a violated preference plus an underlying relevant analogy. Though relevance and metaphor has been discussed, and analogy and metaphor has been discussed, no-one to my knowledge has brought the two discussions together and emphasized the role of relevance in the discovery of an analogy central to a metaphor. The importance of relevance in recognising metaphors has been argued (e.g., Tversky, 1977; Hobbs, 1983). It has also been frequently claimed that the critical match in a metaphorical relation is some correspondence or analogy between two properties (e.g., Wilks, 1978; Ortony, 1979, p. 167; Tourangeau & Sternberg, 1982, pp. 218-221; Gentner, 1983, p. 162). In an analogy, a set of one or more resemblances is found between the particulars of two entities which otherwise are unalike. Consider, for example, an analogy between electric batteries and water reservoirs:

> "The essence of the analogy between batteries and reservoirs is that both store potential energy, release that energy to provide power for systems, etc. We can be quite satisfied with the analogy in spite of the fact that the average battery differs from the average reservoir in shape, color, and substance" (Gentner, 1983, p. 156).

(13) "The ship ploughed the waves."

In (13), which meta5 analyzes, there is a metaphorical semantic relation between 'ship' and the sense of 'plough' meaning "to turn over soil by means of a plough." This sense of the verb 'plough' expects as surface subject the instrument 'plough' (noun). Figure 8 shows the sense-frames in meta5's lexicon for the main (vessel) sense of 'ship', the sense of the verb 'plough' just discussed, together with its preference, the noun sense 'plough'.

```
sf(plough2,
    [[arcs
     [[supertype, transfer1]]],
     [node2,
      [[instrument,
        [preference, plough1]],
       [object,
        [preference, soil1]]]]]).

sf(ship1,                          sf(plough1,
  [[arcs,                            [[arcs,
    [[supertype, watercraft1]]],       [[supertype, tool1]]],
   [node0,                            [node0,
    [[it1, carry1, shipment1],         [[farmer1, plough1, it1],
     [it1, have1, engine1]]]]).         [it1, plough2, soil1]]]]).
```

Figure 8. Sense-frames of ship1, plough2 (verb), and plough1 (noun).

Collation finds a preference violation between plough1 and ship1 because a ship is not a kind of plough. Collation then seeks what is relevant in the sentence; it takes as relevant the action described in (13), which is ploughing. Relevance is then used to highlight any information about ploughs inherited down the sense-network that is especially related to ploughing. In this way, the cell [plough1, plough2, soil1] is selected and dubbed the *relevant* cell. An analogy is then found between this cell and one from 'ship' (shown in Figure 9), hence the metaphorical relation.

Relevant cell of plough1	Cells of ship1
[plough1, plough2, soil1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], [composition1, metal1], [ship1, use2, energy_source1], [boatman1, sail1, ship1], **[ship1, sail2, water2],** [ship1, carry1, shipment1], [ship1, have1, engine1]]

Figure 9. Match of relevant cell from plough1 with cells from ship1.

Relevance is thus used to constrain the semantic interpretation process in CS by picking just one cell from many, a cell which is then given special treatment (an analogy is sought for it). This is admittedly a very limited notion of relevance, where relevance is derived from main sentence verbs and acts on verb preferences. Indeed, the notion of "context" used here is so restricted – relevance operates on parts of the dictionary entry of the word sense which is its source – that this is hardly a contextual constraint at all. However, it is a not unpromising place from which to begin building a more sophisticated theory of relevance: begin by trying to understand the role(s) that individual words (e.g., verbs) play in determining what is relevant and the effect of that notion of relevance on aspects of semantic interpretation (e.g., metaphor recognition). Note for example that the way relevance is used in CS, one lexical semantic constraint (relevance) acts upon a contextual lexical semantic constraint (preferences). This is perhaps a clue to extending the role of relevance: focus upon its effect on contextual lexical semantic constraints.

No notion of salience is presently used in CS but the addition of a simple version is planned. The version would be used to investigate Ortony's (1979) ideas on salience imbalance and metaphor. Ortony argues that in sentences of the form "A is like B,"

> "the imbalance ... in salience levels of matching attributes of the two terms is a principle source of metaphoricity" (Ibid., p. 164).

(14) "Billboards are like placards."

According to Ortony (p. 165), there is likely a literal comparison if the matching attributes of A and B are all of high salience. Hence in (14) there is a literal comparison because billboards and placards share numerous high-salience attributes.

Ortony claims that a nonliteral, metaphorical comparison occurs when (a) no high-salience attributes are shared, (b) some high-salience attributes of B match with less salient attributes of A, and (c) some high-salience attributes of B fail to match with any attributes from B.

(15) "Billboards are like warts."

This is the case in (15) where (a) billboards and warts do not share any high-salience attributes, (b) high-salience attributes of warts, e.g., that they are ugly, match with low-salience attributes of billboards, and (c) other high-salience attributes of warts, e.g., that they are found on the skin, do not match with attributes of billboards (at least they do not *directly* match, I might add).

(16) "Billboards are like pears."

An anomalous comparison occurs, says Ortony, where only low-salience attributes are shared, as in (16).

The simple version of salience to be used in CS, which was pointed out to me by Nicole Modiano, is based on the view that salience is a function of the immediateness of availability of information within a hierarchically organized network (cf. Weiner, 1984, p. 2). The idea is that the most salient information about a word sense is the differentia information available in the node section of its sense-frame. If that section contains a list then the most salient information is at the bottom of the list. Hence in Figure 1, the most salient information about John McEnroe is that he is married, and the next most salient information is that he is male. Decreasingly salient information is found by consulting the superordinates of john_mcenroe and checking the information in their sense-frames, hence the next most salient information about John McEnroe is found in the sense-frame for tennis_player1, which is that he plays tennis. Likewise in Figure 8, the most salient information about plough1 is in its sense-frame. As can be seen, the relevant cell selected is at the bottom of the list and hence is the most salient of the list of cells inherited for plough1, where the most salient information is at the bottom and the least salient is at the top. The same ordering of salience information applies to the list of cells for ship1 shown in Figure 9. Hence [ship1, sail2, water2] is the third most salient cell of ship1 (and [bounds1, distinct1] is the least salient).

This notion of salience acts only on nouns because only they use an inheritance mechanism. To add the notion of salience to CS and meta5 requires modifying that inheritance mechanism to assign numbers to cells inherited down the sense-network indicating their salience ordering. With salience added to CS, there would be two ways to prioritize lexical semantic information. This would be especially useful in metaphor interpretation as it would allow Ortony's (1979) ideas on salience to be explored and should allow further constraints to be imposed on what constitutes an acceptable analogy.

5. Lexical Semantic Structure

A number of studies of lexical structure have made proposals regarding the internal organization of words of various parts of speech. These proposals, I would argue, constitute an investigation of various non-contextual lexical semantic constraints. One set of studies and proposals concern the thematic structure of verbs. Another set, which I will discuss briefly here, concern the assignment of internal semantic structure to nouns. As we saw in Figures 2 and 9 for the sense-frames of John McEnroe and plough1, nouns can be associated with large amounts of information, yet they cannot be structured in terms of preferences and assertions, the two-occurrence constraints which so help organize the semantic information in verbs, adjectives, and other parts of speech.

I will briefly consider four proposals concerning the internal semantic structure of nouns made by Schank and his associates, Dahlgren, Fass, and Pustejovsky. All four assume a naive common sense view of the world, a view expressed by Schank and Abelson (1977) when they declare their focus to be upon

> "the world of psychological and physical events occupying the mental life of ordinary individuals, which can be understood and expressed in ordinary language" (Schank & Abelson, 1977, p.4).

They argue that people use knowledge of "naive psychology,"

> "the common sense (though perhaps wrong) assumptions which people make about the motives and behavior of themselves and others" (Ibid.),

and a kind of "naive physics,"

> "or primitive intuition people have about physical reality" (Ibid.).

Schank and his associates argue that understanding can happen at many levels of description so they have attempted to distinguish some of those levels from the micro-level of naive inferences about "causality" through to a more coarse-grained macro-level of scripts (stereotypical sequences of actions associated with an episodic situation such as going to a restaurant), plans which contain scripts, goals (i.e., motivations and intentions) which give rise to plans, and themes (general beliefs) which affect the goals people pursue (see, e.g., Schank, 1980, pp. 254-257). These levels have been used to characterize nouns such as 'restaurant', 'bus', and 'pickpocket' (see Schank & Abelson, 1977) and give lexical structure to such nouns.

Moreover, Schank and his associates have studied how these lexical structures can constrain semantic interpretation. One constraint explored is restrictions on the kinds of inferences associated with different levels of description (cf. Schank, 1980). Another constraint suggested is an "invariance hierarchy" (Carbonell, 1982, pp. 427-429) for limiting metaphor interpretation. Carbonell suggests that when two terms are compared in a metaphor, as when 'ship' and 'plough' (noun) are compared in (13), certain properties are more likely to be transferred unchanged than others, for example, goals and plans are likely to be transferred intact.

Dahlgren (1988) also assumes a naive theory of the world. She proposes that nouns be structured in terms of 54 "kind types" derived from asking people to list features of objects (see Table 2.5 on p. 62) including colour, material, part of, sex, source, taste, appearance, function, location, physiology, shape, length, and size. These types include perceptual features, functional features, and so on.

> "Features at lower nodes inherit feature type constraints from the nodes above them in the ontology .... At each node, only certain feature types are applicable" (p. 73).

In my own work, sense-frames are organized in terms of a differentia classification of knowledge which is based on seven basic properties that concern fundamental perceptual, spatial, physical, and biological notions – boundedness, extendedness, composition, behaviour, animacy, biological category, and sex – into which even very general notions such as "entity," "physical object," and "abstract object" can be decomposed. For example, a physical object has boundedness, has three-dimensional extent, physical composition, behaves like a solid, and has animacy (it could be inanimate, it could be animate). These properties serve to constrain the definitions supplied to noun senses and simplify matching because any pair of noun sense definitions has most, if not all, of the seven properties.

Pustejovsky (1989) has outlined a theory of "qualia structure" for nouns. Qualia structure consists of four basic roles. The constitutive role comprises the material, weight, and parts and component elements of an object. The formal role includes its orientation, shape, dimensionality, and colour. The telic role includes the use an agent has for the object and the object's function. Finally, the agentive role includes factors responsible for the origin of an object including, for example, its creator and the causal chain that led to its creation.

6. Summary, Discussion and Conclusions

Following Starosta (1988), the following issues about constrainthood were distinguished:

[1] contextual versus non-contextual,

[2] lexical versus non-lexical (e.g., properties of rules and representations),

[3] semantic versus syntactic.

A series of lexical semantic constraints were discussed in terms of their role in linguistic description and semantic interpretation: preferences, assertions, relevance, salience, and theories of the internal semantic structure of nouns. Some of the constraints were contextual, e.g., preferences and assertions, the two co-occurrence constraints. Some were non-contextual, e.g., salience and noun semantic structure. Relevance, the remaining constraint, seemed to lie across the contextual/non-contextual divide.

In the introduction, we observed two main ways in which language is being discussed in terms of constraints. One way is that linguistic constraints are being used to structure some recent descriptions of language research by Grishman and Allen. This is worth pursuing because constraints are a useful way to organize descriptions of syntactic and semantic theories of natural language. For example, they allow us to better understand how some theories of lexical ambiguity resolution work, as in section 2 where preference-based and relational constraint based approaches were discussed.

The second way was that natural language understanding is being increasingly described in terms of the integration of diverse constraints. I would put this down, in part, to the development of some recent approaches to language representation and processing, which are sufficiently general to allow language to be productively thought of in terms of the integration of constraints, notably unification-based approaches (Fenstad et al, 1987; Pollard & Sag, 1987). I should add that architectures amenable to the idea of integrating constraints has been around for quite a while, notably the blackboard model (Erman, Hayes-Roth, Lesser, & Reddy, 1980) which is a global database that records hypotheses generated by "multiple diverse knowledge sources." Each source, in effect, contains a language constraint of some kind (see, e.g., Ibid., p. 219, Figure 2).

Given that some promising architectures exist, we need to describe more precisely how lexical semantic constraints are specified and enforced, to use Grishman's terminology. In particular, the relational constraints mentioned in section 2 could do with more attention. If we could understand to specify and enforce them, we might be able to make some progress in the semantics of prepositional phrase attachment. We also need to understand how different lexical semantic constraints interact, as in the interaction of relevance and preferences in CS, which will become more complicated should salience be introduced. These research issues seem central to lexical semantics; indeed, the study of lexical semantic constraints seems an appropriate focus for lexical semantic research.

7. Acknowledgements

This chapter was prepared while I was on a postdoctoral fellowship supported by the Advanced Systems Institute of British Columbia and the School of Computing Science, Simon Fraser University, BC. I would like to thank Fred Popowich for several helpful discussions about the nature of syntactic constraints and he, Carl Vogel and Veronica Dahl for useful comments on an earlier draft of this chapter.

8. References

Allen, James (1987). *Natural Language Understanding*. Menlo Park, CA: Benjamin/Cummings.

Bar-Hillel, Yehoshua (1960). The Present Status of Automatic Translation of Languages. In Franz Alt (Ed.) *Advances in Computers, Volume 1*. New York, NY: Academic Press, pp. 91-163.

Carbonell, Jaime G. (1982). Metaphor: An Inescapable Phenomenon in Natural Language Comprehension. In Wendy G. Lehnert and Martin H. Ringle (Eds.) *Strategies for Natural Language Processing*. Hillsdale, NJ: Erlbaum Associates, pp. 415-434.

Charniak, Eugene (1981). The Case-Slot Identity Theory. *Cognitive Science*, 5, (3), pp. 285-292.

Chomsky, Noam (1957). *Syntactic Structures*. Janua Linguarum NR 4, The Hague, Holland: Mouton.

Chomsky, Noam (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Cruse, D. A. (1986). *Lexical Semantics*. Cambridge, England: Cambridge University Press.

Dahlgren, Kathleen (1988). *Naive Semantics for Natural Language Understanding*. Norwell, MA: Kluwer Academic Publishers.

Erman, Lee D., Frederick Hayes-Roth, Victor R. Lesser, & D. Raj Reddy (1980). The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty. *Computing Surveys*, 12, (2), pp. 213-253. Reprinted in Robert Englemore & Tony Morgan (Eds.) (1988). *Blackboard Systems*. Wokingham, England: Addison-Wesley, pp. 31-86.

Fass, Dan C. (1983). Preference. In Jock McNaught, Doug Arnold, Dan Fass, Claire Grover, Xiuming Huang, Rod Johnson, Harry Somers, Pete Whitelock, & Yorick Wilks (Eds.). Structure, Strategies, and Taxonomy. Eurotra Contract Report ETL-1-E, European Communities, Luxembourg.

Fass, Dan C. (1988a), Collative Semantics: A Semantics for Natural Language Processing. (PhD thesis) Memorandum MCCS-88-118, Computing Research Laboratory, New Mexico State University, NM.

Fass, Dan C. (1988b). An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution. In Steve L. Small, Gary W. Cottrell, & Michael K. Tanenhaus (Eds.) *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology, and Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann, pp. 151-178.

Fass, Dan C. (1988c). Metonymy and Metaphor: What's the Difference? *Proceedings of the 12th International Conference on Computational Linguistics (COLING-88)*, Budapest, Hungary, pp. 177-181.

Fass, Dan C. (1988d). Collative Semantics: A Study in the Discrimination of Meaning. Technical Report CSS/LCCR TR 88-24, Centre for Systems Science, Simon Fraser University, Burnaby, BC, Canada.

Fass, Dan C., & Yorick A. Wilks (1983). Preference Semantics, Ill-Formedness and Metaphor. *American Journal of Computational Linguistics*, Special Issue on Ill-Formed Input, 9, (3-4), pp. 178-187.

Fenstad, Jens Erik, Per-Kristian Halvorsen, Tore Langholm, & Johan van Bentham (1987). *Situations, Language, and Logic*. Dordrecht, Holland: Reidel.

Fillmore, Charles J. (1968). The Case for Case. In E. Bach & R.T. Harms (Eds.) *Universals in Linguistic Theory*. New York, NY: Holt, Rinehart, and Winston, pp. 1-88.

Gentner, Dedre (1983). Structure Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, pp. 155-170

Grishman, Ralph (1986). *Computational Linguistics*. Cambridge, England: Cambridge University Press.

Gruber, Jerome S. (1965). Studies in Lexical Relations. (PhD thesis), MIT, Cambridge, MA. Published as *Lexical Structures in Syntax and Semantics*. Amsterdam, Holland: North-Holland, 1976.

Hawkinson, L. (1975). The Representation of Concepts in OWL. *Proceedings of the 4th International Joint Conference on Artificial Intelligence (IJCAI-75)*, Tbilisi, Georgia, USSR, pp. 107-114

Hayes, Philip J. (1976). Semantic Markers and Selectional Restrictions. In Eugene Charniak & Yorick A. Wilks (Eds.) *Computational Semantics*. Amsterdam, Netherlands: North-Holland, pp. 41-54.

Hirst, Graeme J. (1987). *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge, England: Cambridge University Press.

Hobbs, Jerry R. (1983). Metaphor Interpretation as Selective Inferencing: Cognitive Processes in Understanding Metaphor (Part 2). *Empirical Studies of the Arts*, 1, pp. 125-141.

Hudson, Richard A. (1984). *Word Grammar*. Oxford, England: Basil Blackwell.

Jackendoff, Ray S. (1976). Towards an Explanatory Semantic Representation. *Linguistic Inquiry*, 7, pp. 89-150.

Kaplan, Ronald, & Joan Bresnan (1982). Lexical-Functional Grammar. In Joan Bresnan (Ed.) *The Mental Representation of Grammatical Relations*. Cambridge, MA: MIT Press, pp. 173-281.

Katz, Jerrold J. (1964). Analyticity and Contradiction in Natural Language. In Jerry A. Fodor & Jerrold J. Katz (Eds.) *The Structure of Language: Readings in the Philosophy of Language*. Englewood Cliffs, NJ: Prentice-Hall, pp. 519-543.

Katz, Jerrold J. (1972). *Semantic Theory*. New York, NY: Harper International Edition.

Katz, Jerrold J., & Jerry A. Fodor (1964). The Structure of A Semantic Theory. In Jerry A. Fodor & Jerrold J. Katz (Eds.) *The Structure of Language: Readings in the Philosophy of Language*. Englewood Cliffs, NJ: Prentice-Hall, pp. 479-518.

Katz, Jerrold J., & Paul Postal (1964). *An Integrated Theory of Linguistic Description*. Cambridge, MA: MIT Press.

Lees, R.B. (1960). The Grammar of English Nominalizations. *International Journal of American Linguistics*, 26, 3.

Lehnert, Wendy G. (1988). Symbolic/Subsymbolic Sentence Analysis: Exploiting the Best of Two Worlds. COINS Technical Report 88-99, Department of Computer and Information Science, University of Massachusetts at Amherst, MA.

Lehrer, Adrienne (1974). *Semantic Fields and Lexical Structure*. Amsterdam, Holland: North-Holland.

Levin, Beth (Ed.) (1985). Lexical Semantics in Review. Lexicon Project Working Papers Number 1, Center for Cognitive Science, MIT, Cambridge, MA.

Olawsky, Duane E. (1989). The Lexical Semantics of Comparative Expressions in a Multi-Level Semantic Processor. *Proceedings of the 27th Annual Meeting of the ACL*, Vancouver, BC, pp. 169-176.

Ortony, Andrew (1979). Beyond Literal Similarity. *Psychological Review*, 86, (3), pp. 161-179.

Ortony, Andrew, Richard J. Vondruska, Mark A. Foss, & Lawrence E. Jones (1985). Salience, Similes, and the Asymmetry of Similarity. *Journal of Memory and Language*, 24, pp. 569-584.

Pollard, Carl (1989). The Syntax-Semantics Interface in a Unification-Based Phrase Structure Grammar. Unpublished ms.

Pollard, Carl, & Ivan A. Sag (1987). *Information-Based Syntax and Semantics. Volume 1: Fundamentals*. Center for the Study of Language and Information (CSLI) lecture notes no. 13, Stanford, CA: CSLI

Pustejovsky, James (1987). On the Acquisition of Lexical Entries: The Perceptual Origin of Thematic Relations. *Proceedings of the 25th Annual Meeting of the ACL*, Stanford, CA, pp. 172-178.

Pustejovsky, James (1989). A Theory of Computational Lexical Semantics. Unpublished ms.

Quillian, M. Ross (1967). Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities. *Behavioral Science*, 12, pp. 410-430. Also reprinted in Ronald J. Brachman & Hector J. Levesque (Eds.) (1985) *Readings in Knowledge Representation*. Los Altos, CA: Morgan Kaufmann, pp. 98-118.

Schank, Roger C. (1975). *Conceptual Information Processing*. Amsterdam, Holland: North-Holland.

Schank, Roger C. (1980). Language and Memory. *Cognitive Science*, 4, pp.243-284.

Schank, Roger C., & Robert P. Abelson (1977). *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Slator, Brian M. (1988). Lexical Semantics, Preference Semantics, and the Machine Analysis of Text. (PhD thesis) Memorandum MCCS-88-143, Computing Research Laboratory, New Mexico State University, NM.

Starosta, Stanley (1988). *The Case for Case: An Outline of Lexicase Grammatical Theory*. London, England: Pinter Publishers.

Testen, David, Veena Mishra, & Joseph Drogo (Eds.) (1984). *Papers from the Parasession on Lexical Semantics*. Chicago Linguistics Society, University of Chicago, IL.

Tourangeau, Roger, & Robert J. Sternberg (1982). Understanding and Appreciating Metaphors. *Cognition*, 11, pp. 203-244.

Tversky, Amos (1977). Features of Similarity. *Psychological Review*, 84, pp. 327-352.

Ullmann, Stephen (1957). *The Principles of Semantics*. Oxford, England: Basil Blackwell & Mott Ltd.

Waltz, David L., & Jordan B. Pollack (1985). Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation. *Cognitive Science*, 9, pp. 51-74.

Weiner, E. Judith (1984). A Knowledge Representation Approach to Understanding Metaphors. *American Journal of Computational Linguistics*, 10, pp. 1-14.

Weischedel, Ralph M., & Lance A. Ramshaw (1987). Reflections on the Knowledge Needed to Process Ill-Formed Language. In Sergei Nirenburg (Ed.) *Machine Translation: Theoretical and*

Methodological Issues. Cambridge, England: Cambridge University Press, pp. 155-167.

Wilks, Yorick A. (1975a). A Preferential Pattern-Seeking Semantics for Natural Language Inference. *Artificial Intelligence*, 6, pp. 53-74.

Wilks, Yorick A. (1975b). Preference Semantics. In Edward L. Keenan (Ed.) *Formal Semantics of Natural Language*. Cambridge, England: Cambridge University Press, pp. 329-348.

Wilks, Yorick A. (1978). Making Preferences More Active. *Artificial Intelligence*, 11, pp. 197-223. Reprinted in Nicholas V. Findler (Ed.) (1979). *Associative Networks: Representation and Use of Knowledge By Computers*. New York, NY: Academic Press, pp. 239-266.

Wilks, Yorick A., Xiuming Huang, & Dan C. Fass (1985). Syntax, Preference and Right Attachment. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI-85)*, Los Angeles, CA, pp. 779-784.

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

•CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.